



What this rests on

An instrumented corridor database, three methods that chain into one another, and what they refuse to claim.

I

The substrate

Before the method, the database. What we have catalogued, where it comes from, and what we derive from it.



We do not write notes off the back of the news. We query a corridor database that we maintain.

Every analysis starts from catalogued, sourced, versioned objects — not from research restarted for each engagement. That is what makes two corridors comparable — the same grid, applied the same way — what makes a change trackable, and what lets us answer tomorrow the question that was asked today.

The order of magnitude

334

CATALOGUED OBJECTS

Straits, canals, ports, land crossings, cables, energy and digital infrastructure.

330

REGISTERED SOURCES

Each carries its reliability level and its licence regime: use, extraction, storage.


15

DERIVATION ENGINES

CVI, control, corroboration, resilience, network, weaponizability, event pressure, exposed trade loss...

Catalogued is not validated: promoting an object to P0 requires sourced evidence and human validation. 30 objects are there today. That figure describes the instructed corpus — the base also serves 2,239 objects, most of which have had no analytical treatment at all.

What “corridor” actually covers



A corridor is not a shipping lane. Three quarters of the corpus are not straits — they are ports, land crossings, cables and energy installations.

Provenance is a constraint, not an appendix

Every source enters the registry with its reliability level and its licence regime. A source whose licence forbids extraction is quarantined by the loader: it never enters the database, and it cannot turn up in a deliverable by accident.

Canonical and derived are separated all the way down into database privileges: the analytics engines run under a role that is read-only on canonical data. An analysis cannot, structurally, rewrite the fact it started from.

The corpus is not frozen: every schema change is a versioned migration, and every architectural decision is written down before it is applied.

DISCIPLINE

Registered sources	330
Derivation engines	15
Written architectural decisions	126
Versioned migrations	92
Pinned read contract	4.0.0

This is not rigour on display: it is what makes a claim refutable six months later. The read interface is itself a versioned contract, pinned on the consumer side: a break is caught at compile time, on both sides, never in production.

What the database computes on its own

- ◇ Control and weaponizability: who holds a node, and how far that power can be instrumentalised.
- ◇ Corroboration: how many independent sources support a claim, and which ones contradict it.
- ◇ Resilience and network: what falls second if the first node gives way, and along which path.
- ◇ Exposed trade loss: the order of magnitude of the flow actually at stake, not the corridor's total volume.
- ◇ Event pressure: the gap between observed activity and its own baseline.

What moves, and how we let it in

- ◇ Per-corridor press aggregation, with media attention spike detection.
- ◇ Raw GDELT feed for worldwide event coverage.
- ◇ Prediction-market consensus, attached to a corridor by an LLM judge, then floored at two markets minimum.
- ◇ A robots.txt gate on all collection: what a publisher refuses, we do not take.

The cardinality floor is recent and it costs: it removed the consensus block from Panama, our most visible corridor. A signal backed by a single market is not a consensus, however convenient.

II

Measure

CVI: eight dimensions to qualify a corridor's vulnerability, and one hard rule about scores.

Eight dimensions, and the question each one asks

01 Exposition

Quels flux dépendent du corridor ?

02 Concentration

Existe-t-il des alternatives crédibles ?

03 Menace

Quels acteurs/événements peuvent perturber ?

04 Capacité de perturbation

Ces acteurs ont-ils les moyens réels de perturber ?

05 Résilience

Combien de temps pour contourner/réparer/absorber ?

06 Coût de contournement

Quel coût économique, logistique, assurantiel ou politique ?

07 Gouvernance

Qui peut sécuriser/coordonner/stabiliser ?

08 Incertitude

Que ne sait-on pas, avec quel niveau de confiance ?

What the scale gives you, tier by tier

- ◇ Public and Basic — a qualitative diagnosis, low → critical: enough to place a corridor and compare it to another.
- ◇ Standard — a 0–5 score per dimension: enough to see WHERE the vulnerability sits, and therefore what to act on.
- ◇ Premium — an aggregate 0–100 score weighted for the client's context. That condition is not met today: until the weighting is documented and published, the aggregate is not computed, and the engine does not serve it.
- ◇ At every tier, each score carries its sources, its date, its confidence level and its uncertainties.
 - No aggregate 0–100 score without published methodological documentation.
 - No probability of disruption: we do not produce them.
 - No navigational or legal precision in map geometry.
 - As of 2026-08-21: the 334 instructed corridors occupy all four bands (127 / 137 / 65 / 5) — six weeks earlier the base discriminated across ONE. It distinguishes better because it asserts less: the concentration dimension was withdrawn from objects with no substitution study, and 275 corridors out of 334 now carry only two dimensions.

III

Surface

HDDE: a guided interview that works back from the declared supplier to the node that actually binds.



The supplier a company declares is almost always the wrong object of analysis.

HDDE starts from the visible actor and works back: what the dependency actually is, which flow is at stake, who can block it, and how fast a rupture reaches the client. The interview is guided because hidden dependency is never volunteered — it is inferred.

The interview, in eleven moves

01 Framing the case

Scope, starting visible actor, decision question.

02 Critical visible actor

The declared supplier, body or point of passage.

03 Nature of the dependency

Contractual, technical, financial, regulatory?

04 Critical flow

The physical, informational or financial flow actually at stake.

05 Real substitution

Is there a credible alternative, at what cost, in what time?

06 Hidden dependencies

Second-tier suppliers, jurisdictions, invisible nodes.

07 Control and disruption

Gatekeepers, regulators, insurers, disruptive actors.

08 Impact timing

How fast a rupture propagates to the client.

09 Decision thresholds

The levels at which to prepare, act or escalate.

10 Red team

Adversarial testing of the diagnosis before any conclusion.

11 Synthesis

A traceable diagnostic packet, source by source.

Nine diagnostic dimensions, scored 0–5

01 Dependency on the visible actor

How far does the company depend on the declared actor?

02 Hidden dependency

What exposure runs through nodes not visible upstream?

03 Substitution weakness

Is replacement credible, fast, sustainable?

04 Jurisdictional exposure

Which jurisdictions can constrain or block the flow?

05 Flow criticality

Does a rupture put the business itself at stake?

06 Time pressure

How long before the impact reaches the client?

07 Gatekeeper pressure

What control or disruption power do intermediaries hold?

08 Evidence quality

Does the diagnosis rest on reliable, validated evidence?

09 Decision readiness

Is the company ready to decide on this exposure?

The eighth dimension scores how solid the other seven are: a diagnosis can be alarming and poorly evidenced, and that has to be visible.

The evidence scale, and what it refuses



The model proposes and challenges; it attests to nothing. An LLM suggestion is never admissible as evidence, however confident it sounds.

What comes out of the interview

- ◇ A versioned diagnostic packet, carrying the hash of the method pack that produced it.
- ◇ A dependency and control matrix: who holds what, and through which lever.
- ◇ A light action layer: what is feasible now, at what cost, to what effect.
- ◇ A diff between two versions: what changed since the last interview, and why.
- ◇ FR and EN exports, plus the canonical JSON packet.

IV

Arbitrate

VERDICT: seven moves, seven weighted criteria, four outcomes — and a mandatory stop condition.

Seven moves, and the guardrail on each

V **Voir la situation réelle**

La situation est posée sans la réponse.

E **Évaluer les forces externes**

Pas de PESTEL encyclopédique.

R **Révéler position, contraintes et asymétries**
Une force sans preuve devient une hypothèse.

D **Définir les options décisionnelles**
Alternative minimale et opposée/non-action obligatoires.

I **Interroger hypothèses, preuves et biais**

Preuve et contre-argument explicités.

C **Comparer par score, risques et vetos**

Le score n'est pas la décision ; les vetos sont vérifiés.

T **Trancher, tester ou différer**
Condition d'arrêt obligatoire.

Seven criteria, weighted out of one hundred

Valeur stratégique		20	L'option améliore-t-elle réellement la situation ?
Adéquation au contexte		15	Résiste-t-elle aux forces externes ?
Capacité réelle		15	Avons-nous vraiment les moyens d'exécuter ?
Viabilité systémique		15	Peut-elle fonctionner comme un système réel ?
Risque net		15	Le risque est-il acceptable au regard du gain ?
Niveau de preuve		10	Repose-t-elle sur des preuves solides ?
Optionalité		10	Ouvre-t-elle plus de portes qu'elle n'en ferme ?

The weighting is shown before the assessment, not tuned after it. A weight discovered at the end is a result dressed up as a method.

Four outcomes, and the score band that leads there

≥ 80 DO

Commit. Evidence ≥ 4 , no veto, human validation.

60–79 TEST

Run a falsifiable, bounded test, with its stop condition.

40–59 DEFER

Evidence is missing or the context is not ready. Set the review date.

0–39 ABANDON

The option does not hold. Write it down, so it is not revived in six months.

What overrides the score

- ◇ Hard vetoes, audited separately: a high score does not survive a regulatory, financial or capacity impossibility.
 - ◇ Explicit, bounded adjustments, applied as reasoned penalties, never as rounding.
 - ◇ At least three options, including a minimal alternative and the opposite or doing nothing.
 - ◇ A mandatory stop condition: how you will know the decision was wrong, and on what date you look.
-
- No option is selected on score alone.
 - No conclusion is issued without an adversarial red-team pass.
 - No decision is taken by the model: validation is human and nominative.



The score is not the decision. It is what the decision will have to justify itself against.

VERDICT does not produce an automatic arbitration: it produces a reasoned one, with its assumptions, its evidence, its contradictions and its review date. Responsibility stays whole, and that is the point.

V

Walk it through

The method applied to Malacca, on public sources. An illustration, not a published diagnosis.

The corridor, and what actually moves through it

Malacca is not a route, it is a narrowing. The binding passage is the Phillips Channel, a few kilometres off Singapore — and that is where everything that follows concentrates.

The first task is not to estimate a risk, it is to establish what passes. Crude volume, the share bound for a single importer, transit counts, the adjacent hub’s container throughput: four institutional measures, all sourced, none derived.

Without this step, everything after it would be an opinion about a geography.

MEASURED, SOURCED	
Crude in transit	23.2 Mb/d
Share imported by China	48%
Transits (Malacca + Singapore)	94,301 / yr
Hub container throughput	41.1 M TEU

Where the dependency hides — and what nobody documents

The naive question is “can ships go elsewhere?”. The answer is yes: Lombok is deep and permissive, it takes vessels Malacca turns away. The constraint is therefore not draught.

The constraint is the fleet. A detour of 300 to 1,000 nautical miles cuts annual rotations and ties up more vessels: rerouting hits systemic absorption, not geography. The only quantified alternative, the China–Myanmar pipeline, covers roughly 2% of the strait’s throughput.

And the most useful result of the analysis is an absence: no public source quantifies residual absorption capacity. We write that down.

SUBSTITUTION, EXAMINED

Lombok depth	≈ 1,000 m
Route extension	+1 to +4 days
China–Myanmar pipeline	≈ 400 kb/d
Share of throughput covered	≈ 2%
Residual capacity	undocumented

Method illustration on public sources — this is not a published diagnosis. The corresponding Atlas entry is still in human validation. Sources: Reuters (2017) and Global Energy Monitor for pipeline capacity; the route extension is a model, not official data.

The threshold, and the decision it triggers

A diagnosis that does not say what to watch tomorrow morning has served no purpose. Each threshold ties an observable indicator to a regime shift and to the action it implies.

They are explicitly graded: backed by verified sources, partial, or hypothetical marker. A hypothetical threshold is still useful — provided it does not pass itself off as a measurement.

Crossings matter less than co-occurrence: a volume drop alongside rising incidents signals a durable shift, where either signal alone stays cyclical.

DECISION MARKERS

Crude flow decline	> 15%
Transit decline (4 wks)	> 20%
Sustained SOMS incidents	≥ 1 / wk × 3
Continuous physical blockage	> 72 h
SOMS incidents 2025	108 (+74%)

Method illustration on public sources — this is not a published diagnosis. The corresponding Atlas entry is still in human validation. Source: ReCAAP ISC, Annual Report 2025. Thresholds are decision markers — an analysis, never a measurement.

The threshold is crossed. What do we do, and how will we know?

A crossed threshold is not a decision, it is an entry into the protocol. The situation is stated without its answer, at least three options are defined — the main one, a minimal alternative, and doing nothing — and only then are they compared.

In this case, the option “pre-position thirty days of stock on inputs transiting Malacca” comes out with partial evidence: the flows are measured, the residual absorption capacity is not. The score reflects that, deliberately — the weighting gives ten points to the level of proof.

So the outcome is not to commit. It is a bounded test whose stop condition is written before it starts: what will tell us the decision was wrong, and on what date we look.

THE ARBITRATION, AS IT COMES OUT

Triggering threshold	-15 % crude
Options compared, incl. non-action	3
Level of proof retained	3 / 5
Hard vetoes	none
Outcome	60-79 · TEST

Going further

The full method, corridor by corridor, is run inside a closed six-to-eight-week Premium pilot.

The framing interview precedes any proposal — it exists as much to qualify the request as to decline it if it is not ours to take.

contact@applied-geopolitics.com

www.applied-geopolitics.com